

[Claim 41] (original) A system according to claim 1 will notify the producer of the node distribution message of success or failure.

[Claim 42] (original) Distributed nodes according to claim 14 may be distributed at deployment time.

[Claim 43] (original) A system according to claim 1 will set an attribute of the input node to indicate that the pattern has been distributed.

[Claim 44] (original) An input node according to claim 43 will maintain a list of the remote systems and the total number of data objects routed remotely.

[Claim 45] (original) A system according to claim 1 may not attempt to distribute nodes that were distributed by another rule engine. Instead, it should notify the originating engine it cannot receive additional data until resources are available.

[Claim 46] (original) A system according to claim 1 may randomly select a remote engine to route data to, if the pattern is distributed to more than 1 engine.

ABSTRACT

[Para 6] Distributed pattern matching defines novel extensions and improvements to the RETE algorithm and to the forward chaining method of inference in rule systems. Prior implementations of Dr.Forgy's RETE algorithm do not distribute reasoning. This includes enhancements to RETE like RETE-UL, RETE*, TREAT and LEAPS algorithm. All prior research in the field of pattern matching algorithms address the efficiency within a single system. Existing research in the field of distributed reasoning did not consider the techniques described in the patent claims. Due to the increased performance of personal computers and gigabit Ethernet network, it is feasible to distribute reasoning across several systems. The same approach is envisaged for a goal-driven, backward chaining inference method where data is used to satisfy new goals in an interactive session between different engines, different systems, or between a system and a user interface. Utilizing techniques from Remote Procedure Calling (RPC), distributed processing, parallel processing and grid

computing, the invention defines a novel algorithm to implement a system capable of distributed pattern matching. Reasoning over large data sets exceeding millions of rows in a reactive mode with a single engine is difficult due to hardware limitations. Using a cluster of inexpensive systems, reasoning over large datasets is more efficient and provides better scalability and response time. Distributed pattern matching is a significant innovation and enables rule engines to reason over large sets of data. Current rule engines scale vertically by upgrading the hardware. Distributed pattern matching has no preference towards vertical or horizontal scaling. There are classes of problems that require a rule engine to reason over the entire database. This type of problem is difficult to solve with collaborative reasoning, where multiple symmetric systems are used to process discrete blocks of data. Distributed pattern matching overcomes the limitations of collaborative reasoning and helps solve a new class of problems. This innovation will also facilitate system integration, lower costs of maintenance in large organizations, and improve the performance and quality of applications in many areas, including: electronic procurement, e-government, real-time data management systems, insurance systems, enterprise resource management, and personalization, to just cite a few applications.